

Remarks

5 The following *Remarks* will discuss Examiner's objections and rejections in the order in which they are presented in the Office action. In discussing the prior art, they will deal not only with the Draper reference, but also with IBM Web Cache Manager reference (henceforth "IBM") cited by Examiner in the International Search Report which Examiner made in PCT/US00/10516, which claims priority from the present application. The IBM reference was
10 supplied to Examiner in an IDS mailed 8/1/00.

The amendments to the Abstract, Specification, and claims

The amendment of the *Abstract* shortens it to less than 250 words and also removes the words
15 "the preferred embodiment". The amendment to the Specification fixes typographical errors noted by Applicants. Neither amendment adds new matter. Examiner will also easily ascertain that the amendments to the claims add no new matter. Claims 24, 29-32 have been amended to solve problems with the language noted by Applicant, specifically the lack of an antecedent for "application program" in claim 24 and the use of the non-word "likeliness" in 29-32.

New claims 36-75

20 Claims 36-75 have been added to better claim Applicants' invention. Claim 36 is a method claim corresponding to apparatus claim 1; claims 37-39 are method claims corresponding to claims 2-4; claim 40 is a method claim corresponding to claim 5; claims 41-47 are new method
25 claims dependent from method claim 20; claim 48 is a new apparatus claim dependent from claim 21; claims 49-52 are method claims corresponding to apparatus claims 21-23; claims 53-55 are method claims dependent from method claim 35; claims 56-59 are method claims dependent from either of claims 27 or 28; claim 60 is a method claim dependent from claim 34. Claim 61 is a "carrier" claim corresponding to claim 36; claim 62 is a carrier claim
30 corresponding to claim 37; claim 63 is a carrier claim corresponding to claim 40.

Claims 64-75 are generic claims to Applicants' invention in its aspect as a network server that serves a data object in response to a network handle that is defined by the network and that has

a cache for other objects that are not specified by network handles. Claims 64-69 are generic apparatus claims; claims 70-75 are corresponding method claims. Examiner will easily determine from the claims dependent on claims 64 and 75 that these claims are indeed generic to the Web server disclosed in the present patent application and consequently introduce no new
5 matter.

The rejections under 35 U.S.C. 112, second paragraph

In claims 2-5, 10, 14-19, Applicants have replaced “a probable future request” with language
10 like that of amended claim 2:

when the information fetcher determines that a future request for the information
is probable

Applicants believe that this language should make what is claimed clear, particularly when it is
15 read in conjunction with page 9, line 31-page 10, line 23, which also give examples of two species of the genus claimed in claims 2 and 5. The species themselves are claimed in claims 3 and 4. The objectionable language also appears in other contexts in claims 10, 14, 15, and 19, with claims 11-13 and 16-19 addressing various techniques for “determin[ing] that a future request for the information is probable”. While Examiner does not reject these claims on any
20 basis other than 35 U.S.C. 112, second paragraph, it should be pointed out here that Draper discloses nothing whatever about how an item is selected for encachement, while the IBM reference discloses a straightforward cache on miss technique. See IBM, page 3, third paragraph.

The rejections under 35 U.S.C. 102(e)

Claims 1, 6-9, 20-35 can be divided into two groups by their content: in the first group, the independent claims 1, 6, 9, 20 are distinguished from the Draper reference and the IBM reference by the manner in which the cache is updated; in the second group, the independent claims 21, 24, and 34 are distinguished from Draper and the IBM reference by the interaction
30 between the data access layer and a queryable cache.

The rejection of claims 1, 6, and 20 and claims dependent therefrom

These claims address the manner in which the “partial copy” in the server is updated from the “information source”. Updating in the server is done by “an updater” that is described in claim 1 as follows:

5 an updater that updates information in the partial copy [in response to] from
 update information which the information source provides the server, the update
 information being provided without intervention by the server when information
 in the partial copy is updated in the information source.

As set forth in claim 1, the distinction between the way this updater works and the way
 10 updating is done in Draper or the IBM reference is that the *information source* provides the
 information *without intervention by the server* when information in the partial copy is updated
 in the information source. In Draper, as pointed out at col. 13, lines 41-48, an update operation
 is *always* initiated by the cache. The master database *never* provides an update to the cache on
 its own as a consequence of a change in the item for which the update is provided. In the IBM
 15 reference, synchronization is discussed at the bottom of page 4. There are three options:
 following the “time to live” field in the HTML object, deleting the object after an elapsed time
 from modification, and checking validity on every reference. Thus, though the techniques are
 different, the IBM reference, like Draper, puts responsibility for deciding when to update an
 object in the cache on the cache, not on the information source, as claimed in independent
 20 claims 1,6, and 20. Since there is nothing in Draper or IBM either singly or in combination that
 corresponds to Applicants’ updater, Applicants’ independent claims 1, 6 and 20 are patentable
 over the reference as are all claims dependent from those claims.

As for the claims that are dependent from claim 6, claim 7 is to the trigger mechanism which
 25 causes the information source in a preferred embodiment to provide the updates. There is no
 discussion at all of trigger mechanisms in the IBM reference. Examiner cites col. 5, lines 19-46
 of Draper as disclosing triggers, but the cited location only describes Draper’s tags. As may be
 seen by comparing the discussion of the updating mechanism at page 16, lines 8-31 of
 Applicants’ Specification with the cited location in Draper, Draper’s tags have nothing
 30 whatever to do with Applicants’ triggers. Thus, claim 7 is patentable both because it is
 dependent from claim 6 and in its own right.

Claims 10-14 are further patentable in their own rights for the reasons set forth in the discussion of the rejection under 35 U.S.C. 112, second paragraph set forth above.

The rejection of claims 21, 24, and 34

5 These independent claims are addressed to two features of Applicants' systems: that the cache is *queryable* and that it is the *data access layer* which determines whether to query the cache or the remote database. The two features together make the cache *transparent* to application programs that run on the server and make queries to remote databases. As set forth at page 5, lines 15-19, data is in queryable form when it is in a form to which a query can be applied. As
10 further set forth at page 8, lines 21-27, in Applicants' cache, the cached data is contained in a queryable database, and if the data is in the cache, it can be accessed by a query on the cache. Prior-art data access layers are described in overview at page 2, lines 16-24 of Applicants' specification and the data access layer of Applicants' system is contrasted with prior-art data access layers at page 7, line 32-page 8, line 13.

15

Claim 24 is exemplary for the manner in which these two features are set forth in Applicants' independent claims 21, 24, and 34:

24. (amended) An improved server of the type that provides a program with a standard interface for querying remote datasets,
the improvement comprising:
5 a queryable cache that contains copies of certain of the datasets and is local to the server,
the improved server receiving a query for a remote dataset in a form required by the interface from the program, determining whether a copy of a dataset to be queried is present in the queryable cache, and, if the copy is present, querying the
10 copy, and otherwise querying the remote dataset,
whereby the queryable cache is transparent to the program.

In rejecting these claims, Examiner points out, accurately enough, that Draper discloses first, that the distributed database that employs his tags may be any kind of database, including a relational, and therefore queryable, database, and second, that Draper's
15 system may be used to cache documents that have been converted to HTML pages. Applicants can also add that the IBM reference also discloses a system that caches

HTML pages and further discloses at page 3, third paragraph, a Web proxy server that “intercepts Internet requests from users and routes them to the appropriate storage location”, be it the cache or the host. Examiner also indicates that claims 21, 24, and 34 are not patentable for the same reasons as claim 1, but these claims simply do not address
5 how the queryable cache is updated and therefore are related to claim 1 only in that both involve caches.

With regard to Draper, what is being addressed in claims 21, 24, and 34 is the manner in which Applicants’ server interacts with Applicants’ queryable cache. There is no
10 disclosure at all concerning such an interaction in Draper. Examiner refers Applicants to col. 9, lines 33-53, but the cited location discloses only that one embodiment of Draper’s system is used to cache HTML pages made from documents. These HTML pages do not contain the results of queries made by the server, and consequently, there is no need in the described embodiment for an interaction like that set forth in claim 21. See in that
15 regard page 5, lines 13-21 of Applicants’ Specification. There is no disclosure in Draper whatever of how a program might interact with a version of Draper’s system that had a distributed relational database, for example.

As for the IBM reference, the system as described in the reference caches only Web
20 pages and objects (page 2, first paragraph) and is thus an example of a system in which data contained in the Web pages and objects is not cached in queryable form, as described at page 5, lines 13-21 in Applicants’ Specification. As with Draper’s exemplary system, there is no need in such a system for the kind of interaction between the server and the cache set forth in Applicants’ claims 21, 24, and 34. These claims are
25 thus patentable over either Draper and the IBM reference individually or the combination of the two references, as are all of the claims dependent from these claims.

Turning now to the claims that are dependent from claims 21, 24, and 34, claims 22, 23 ,
25, and 35 are all addressed to ramifications of the fact that in Applicants’ system, global
30 identifiers are used to identify datasets in the remote data sets and local identifiers are used for the copies in the cache. The IBM reference does not disclose anything like the

global and local identifiers and their interactions with the query analyzer, and there is indeed no need for such arrangements in the IBM system, since what it caches are HTML pages, which are always referred to by their URLs. Examiner cites col. 1, line 56, to column 2, line 6, col. 5, lines 47-65, and col. 6, lines 8-48 of Draper as showing global
5 and local identifiers, but the first location indicates only that data base objects *have* identifiers, the second discusses Draper's tags, which contain event identifiers, but are not themselves identifiers of datasets, and the third discusses operations using tags. Again, even if tags were taken to be object identifiers, there is nothing disclosed in Draper which corresponds in any way whatever to the query analyzer of these claims, and
10 these claims are thus patentable in their own rights over the references.

Claims 27-33 address the techniques used in Applicants' system for deciding what datasets are to be added to Applicants' cache and are patentable in their own rights for the same reasons that claims 2-5, 10, 14-19 are patentable.

Patentability of the new claims

New claims 64-75 distinguish Applicants' network server from other network servers with caches on the basis of the fact that Applicants' network server provides first objects in response to network handles and has a cache for second objects that contain
20 information that is included in the first objects, but responds to specifiers for the information that are not network handles. That characteristic of the cache immediately distinguishes it from servers such as Web servers with caches for Web pages. Both the IBM reference and Draper disclose caches of Web pages, but neither discloses a server that provides first objects in response to network handles but caches second objects that
25 are not referenced by the network handles.

Conclusion

Applicants have amended their *Abstract* to overcome Examiner's objections thereto, have amended their Specification and claims to deal with typographical and other linguistic
30 errors discovered therein by Applicants and to simplify their claims, have further amended claims 2-5, 10, 14-19 to overcome Examiner's rejection under 35 U.S.C. 112,

second paragraph, have amended claims 1-6, 9-10, 14-20, and 33 to overcome the rejection as anticipated by Draper, and traversed the rejection of claims 21-35. Applicants have further added new claims 36-75 and shown that they are patentable. As pointed out above, none of the amendments has added new matter. Applicants have consequently complied with the requirements of 37 C.F.R. 1.111(b) and respectfully request that Examiner continue with the examination and allow the claims as amended, as provided by 37 C.F.R. 1.111(a).

Respectfully submitted,

Gordon E. Nelson

Attorney of record,
Gordon E. Nelson
57 Central St., P.O. Box 782
Rowley, MA, 01969,
Registration number 30,093
Voice: (978) 948-7632
Fax: (617) 788-0932

10/20/00
Date

Certificate of Mailing

I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to:

Assistant Commissioner for Patents
Washington, D.C. 20231

on 10/20/00
(Date)

Gordon E. Nelson, #30,093

Gordon E. Nelson
(Signature)